arXiv:1512.03077v1 [cs.SY] 8 Dec 2015

A Survey of Code Optimization Methods for Kalman Filter Extensions

Matti Raitoharju
Robert Piché

December 11, 2015

Abstract

Kalman filter extensions are commonly used algorithms for nonlinear state estimation in time series. The structure of the state and measurement models in the estimation problem can be exploited to reduce the computational demand of the algorithms. We review algorithms that use different forms of structure and show how they can be combined. We show also that the exploitation of the structure of the problem can lead to improved accuracy of the estimates while reducing the computational load.

1 Introduction

Bayesian estimation is used to infer probability distribution of a state in time series. Bayesian estimation works by updating a prior distribution with a measurement to obtain a posterior distribution. The posterior distribution is propagated in time to become a new prior distribution. Bayesian estimation has a wide range of applications from positioning [Müller et al. 2014] and tracking [Bar-Shalom et al. 2002] to brain imaging [Hiltunen et al. 2011] and modeling spread of infectious diseases [Keeling and Rohani 2008]. In general the Bayesian estimate cannot be computed in closed form. Under certain conditions closed form solutions exist. For example, when measurement and state transition functions are linear and associated noises Gaussian, the algorithm known as the Kalman Filter (KF) produces optimal Bayesian estimates [Ho and Lee 1964].

When the conditions are not met an approximation has to be used. Kalman Filter Extensions (KFEs) are based on approximate linear-Gaussian models. In literature, there are several different types of KFEs, with different demands on computational resources. The computational complexity of the KF increases when the number of estimated state variables or dimensionality of measurement vector increases; in some KFEs even exponentially [Ito and Xiong 2000]. In many applications the computational resources are severely limited e.g. when doing tracking using wearable devices.

A drawback of KFEs is that the estimate is approximated as a Gaussian, which makes the estimate inaccurate when the true posterior is not normal. Gaussian Mixture Filters (GMFs) (a.k.a. Gaussian sum filters) use a sum of normal densities to estimate the probability distributions and can approximate any probability density function [Sorenson and Alspach 1971]. Because GMFs use several KFEs

This work is supported by the Academy of Finland.
Author's addresses: Tampere University of Technology Department of Automation Science and Engineering P.O. Box 692, FI-33101 Tampere, Finland. first.last@tut.fi

the computational load is larger than with algorithms that use only one Gaussian. The optimizations in this survey can be applied also in the GMFs for state propagation and update of individual components.

In this survey, we study different methods to optimize the state estimation with KFEs using the structure of the state transition and measurement models. The presented algorithms are such that when applied to situation where the KF produces an optimal result, the result is still optimal. We give the algorithms in a general form so they can be applied to as wide range of problems as possible, but still in a form that they are easy to implement. In addition to surveying the algorithms in literature, we give some generalizations of the algorithms and new ways of applying the optimizations with the KFEs.

For implementations we assume that a library for basic linear algebra is available. In many practical applications the algorithms can be optimized further by taking the application specific structure of the matrices into account. Often the matrices are sparse, and the sparsity can be exploited for optimization either by hand or by implementing the algorithms with sparsity-optimized subroutines. These can be found, for example, in Sparse Basic Linear Algebra Subprograms or in Matlab.

Rest of this survey is organized as follows: In the next section we present the common notations that are used throughout the paper. In Section 3 the background of KFEs is presented. Section 4 contains linearity based optimizations of the KFEs. Section 5 presents optimizations based on unobserved and static state variables. In Section 6 solving the set of linear equations containing the inverse of the innovation covariance is optimized. Section 7 gives example applications of the optimization methods. Section 8 concludes the article.

2 Notations

In the following list there are the variables that are used throughout the paper. In different sections there are algorithm specific notations that are explained as they occur.

- Scalars

  j number of iterations

  m dimension of measurement

  n dimension of state (or augmented state)

- Random variables

  x state

  z augmented state

  ε measurement noise

  ε^x state transition noise

  ε^y measurement noise

- Subscripts

`http://math.nist.gov/spblas/`
`http://www.mathworks.com/products/matlab/`

i iteration index

t time index

- Functions

  $f(\cdot)$ state transition function

  $g(\cdot)$ non-specific function

  $G(\cdot)$ Matrix valued function

  $h(\cdot)$ measurement function

- Expected values

  $\mathrm{E}\, x$ Expected value of x

  $P_{g(x)x}$ Covariance between $g(x)$ and x

  μ_x Mean of x

- Other variables

  I identity matrix

  J matrix that defines statistically linearized relationship between state and measurement, which is the Jacobian matrix in linear systems

  K Kalman gain

  S innovation covariance

  $\mathbf{0}$ zero matrix

- Acronyms

  BS Basestation

  CKF Cubature Kalman Filter

  DD Divided Difference

  EKF Extended Kalman Filter

  EKF2 Second Order Extended Kalman Filter

  GF Gaussian Filter

  GMF Gaussian Mixture Filter

  KF Kalman Filter

  KFE Kalman Filter Extension

  PDR Pedestrian Dead Reckoning

  QKF Gauss-Hermite Quadrature Kalman filter

  RBPF Rao-Blackwellized Particle Filter

  RUF Recursive Update Filter

  SLAM Simultaneous Localization and Mapping

  SOKF2 Second Order Polynomial Kalman Filter

  S^{2}KF Smart Sampling Kalman Filter

  TDoA Time Difference of Arrival

  UKF Unscented Kalman Filter

  URUF Unscented Recursive Update Filter

3 Background

In discrete-time Bayesian filtering, the state of a dynamical system is estimated based on noisy measurements. The state model describes how the n-dimensional state x is propagated in time and can be written as

$$x_t = f_t(x_{t-1}, \varepsilon_t^x),$$

where $f(\cdot)$ is the state transition function, x_{t-1} is the state of the previous time and ε_t^x is the state transition noise. To shorten notations we use augmented state $z = \begin{bmatrix} x^T & \varepsilon^T \end{bmatrix}^T$, where applicable also we do not show the time index t when its absence should not cause confusion. The state's distribution is updated using measurements of form

$$y_t = h_t(z_t),$$

where y_t is the realized measurement, $h(\cdot)$ is the measurement function and z_t is the augmented state of the state and measurement error ε^y.

In general, a KFE can be divided into two parts:

1. Prediction:

$$\begin{aligned} \mu_t^- &= \mu_{f(z_{t-1})} \\ P_t^- &= P_{f(z_{t-1})f(z_{t-1})}, \end{aligned}$$

   where μ_t^- is the prior mean computed using the state transition function and posterior of the previous time step and P_t^- is the prior covariance.

2. Update:

$$\begin{aligned} y_t^- &= \mu_{h(z_t^-)} \\ S_t &= P_{h(z_t^-)h(z_t^-)} && (1) \\ K_t &= P_{x_t^- h(z_t^-)} S_t^{-1} && (2) \\ \mu_t &= \mu_t^- + K(y_t - y_t^-) \\ P_t &= P_t^- - K_t S_t K_t^T, \end{aligned}$$

   where y_t^- is the predicted mean, $z_t^- \sim \mathrm{N}(\mu_t^-, P_t^-)$, S_t is the innovation covariance, K_t is the Kalman gain, μ_t is the updated state, and P_t is the updated covariance.

The mean and covariance matrices associated to measurement and state transition functions can be written using expectations:

$$\mu_{g(z)} = \mathrm{E}\, g(z) \tag{3}$$

$$P_{g(z)g(z)} = \mathrm{E}(g(z) - \mu_{g(z)})(g(z) - \mu_{g(z)})^T \tag{4}$$

$$P_{zg(z)} = \mathrm{E}(z - \mu_z)\left(g(z) - \mu_{g(z)}\right)^T. \tag{5}$$

In (2) $P_{x_t^- h(z_t^-)}$ refers to the rows that correspond to the state variables of $P_{z_t^- h(z_t^-)}$. When a function is linear

$$g(z) = Jz,$$

the expectations have analytic form:

$$\mu_{g(z)} = J\mu_z$$
$$P_{g(z)g(z)} = JP_{zz}J^T$$
$$P_{zg(z)} = P_{zz}J^T$$

and the algorithm is the KF.

There are extensions that approximate the expectations (3)-(5) using analytic differentiation (or integration) of the functions, for example the Extended Kalman Filter (EKF) uses the first order Taylor expansion and Second Order Extended Kalman Filter (EKF2) does the linearization based of the second order Taylor expansion [Jazwinski 1970]. There are also algorithms that make numerical differentiation to approximate (3)-(5). In Divided Difference (DD) filters [Nørgaard et al. 2000] the computation is based on numerical differentiation of the functions to obtain a linear model. DD filters use $2n + 1$ function evaluations. Second Order Polynomial Kalman Filter (SOKF2) [Ito and Xiong 2000] uses $\frac{n^2}{2} + \frac{3}{2}n + 1$ sigma-points to fit a second order polynomial to the function and then compute the analytic covariance matrices for the polynomial.

One very commonly used and large group of KFEs is algorithms that can be written in form:

$$\mu_{g(z)} \approx \sum w_i^s g(\chi_i) \tag{6}$$
$$P_{g(z)g(z)} \approx \sum w_i^c (g(\chi_i) - \mu_g(z))(g(\chi_i) - \mu_g(z))^T \tag{7}$$
$$P_{zg(z)} \approx \sum w_i^c (\chi_i - z)(g(\chi_i) - \mu_g(z))^T, \tag{8}$$

where χ_i are so-called sigma-points that are chosen according to the prior distribution and w_i^s and w_i^c are associated weights. Examples of this kind of filters are Unscented Kalman Filter (UKF) [Wan and Van der Merwe 2000], different Cubature Kalman Filters (CKFs) [Arasaratnam and Haykin 2009], and Gauss-Hermite Quadrature Kalman filter (QKF) [Ito and Xiong 2000].

The selection and amount of sigma-points and weights depend on the algorithm used. The UKF is used usually with $2n + 1$ sigma-points. The Gaussian Filter (GF) uses $kn + 1$ sigma-points, where $k > 1$ is an integer parameter. CKFs are developed for different orders and they use $\mathcal{O}(n^m)$ sigma-points, where m is the order. The number of sigma-points in QKFs increases exponentially as the number of sigma-points is α^n, where $\alpha > 1$ [Ito and Xiong 2000]. There are also algorithms that allow arbitrary number of points, for example Smart Sampling Kalman Filter (S^2KF) [Steinbring and Hanebeck 2014], which uses at least $n + 1$ sigma-points.

Algorithms that use (6)-(7) do not compute the J explicitly. In some optimizations the J is required and can be computed using statistical linearization [Särkkä 2013]

$$J = P_{g(z)z}P_{zz}^{-1}. \tag{9}$$

This matrix defines the relationship between P_{zz} and $P_{g(z)z}$. Because this computation requires solving a set of linear equations it should be avoided when the dimension of z is large.

The computational requirements of the algorithms can be reduced in several ways. One way is to compute the expected values numerically only for state variables that are transformed by a nonlinear function and compute the expectations

for linear parts analytically. The improving implementation efficiency using the linear KF update for the linear state variables is more familiar in particle filtering. Specifically Rao-Blackwellized Particle Filter (RBPF) solves the estimates of conditionally linear variables using a KF and rest of the variables using particles [Doucet et al. 2000]. We can also optimize the dimension of the output of the nonlinear part. In (7) the $m \times m$ covariance matrix is updated for every sigma point. Thus, halving the dimension m of $g(\cdot)$ reduces the operations applied to the covariance matrix to one fourth. This kind of update optimizations are considered in Section 4.

When some state variables are unobservable i.e. they do not contribute to the output of the measurement function and are static i.e. their mean and covariance does not change, they do not need to be updated. Computing their estimate only when they are observed or their mean or covariance changes can also be used to reduce the computational burden. Section 5. It is also possible to compute the Kalman gain (2) faster using the structure of the innovation covariance (1) as shown in Section 6.

4 Optimizations Based on the Linearity in Nonlinear functions

4.1 Partially linear functions

In [Briers et al. 2003] algorithms for treating different setups of UKF were presented. The algorithms considered different implementations when the state transition model or measurement model are of the form

$$\begin{aligned} g(x, \varepsilon) &= Gx + \varepsilon \\ g(x, \varepsilon) &= g_1(x) + \varepsilon \\ g(x, \varepsilon) &= g_2(x, \varepsilon). \end{aligned}$$

The algorithms improve the computation efficiency when either or both of state and measurement models belong to the first or second group. If both are in the first group the updates can be computed optimally using the KF. If the function belongs to the second group then (3)-(5) become

$$\mu_{g(z)} = [\mathrm{E}\, g(x)] + \mu_\varepsilon \tag{10}$$

$$P_{g(z)g(z)} = \left[\mathrm{E}(g(x) - \mu_{g(x)})(g(x) - \mu_{g(x)})^T\right] + P_{\varepsilon\varepsilon}$$

$$P_{xg(z)} = \mathrm{E}(x - \mu_x)\left(g(x) - \mu_{g(x)}\right)^T \tag{11}$$

and the approximation can be computed for the dimension of the state instead of the dimension of the augmented state. This optimization is widely used among different filters and often considered to be the standard model.

In [Morelande and Ristic 2006] the function is split into nonlinear $g_1(\cdot)$ part that depends only on a part of the state z_n and linear parts

$$h(z) = \begin{bmatrix} g_1(z_n) \\ H_1 z \end{bmatrix}.$$

The UKF is used for computing the expected values of the nonlinear part and the correlation between nonlinear and linear parts.

We generalize the above models to functions of form:

$$g(z) = Ag_n(Tz) + Hz, \tag{12}$$

where $\operatorname{rank} T$ equals the number of rows in T. To reduce the required resources in computation T is chosen to have minimal number of rows and $g(\cdot)$ to have minimal number of elements. Expectations (3)-(5) are computed for Tz, which should have a smaller dimension than z.

For computing the update, an algorithm that approximates expectations (3)-(5) is required. These are usually computed in the update stage of a KFE. The cross correlation matrix $P_{zg(z)}$ is not needed in the state propagation, but the algorithm from the update stage can be used for this.

The transformed augmented state is denoted

$$\tilde{z} = Tz \sim \mathrm{N}(\mu_{\tilde{z}} = T\mu_z, P_{\tilde{z}\tilde{z}} = TP_{zz}T^T).$$

When expectations $\mu_{g_n(\tilde{z})}$ (3),$P_{g_n(\tilde{z})g_n(\tilde{z})}$ (4), and $P_{\tilde{z}g_n(\tilde{z})}$ (5) are known, the expectations for (12) are

$$\begin{aligned}
\mu_{g(z)} &= A\mu_{g_N(Tz)} + H\mu_z \\
P_{g(z)g(z)} &= \begin{bmatrix} A & H \end{bmatrix} \begin{bmatrix} P_{g_n(\tilde{z})g_n(\tilde{z})} & P_{zg(\tilde{z})}^T \\ P_{zg_n(\tilde{z})} & P_{zz} \end{bmatrix} \begin{bmatrix} A^T \\ H^T \end{bmatrix} \\
P_{zg(z)} &= \begin{bmatrix} P_{zg_n(\tilde{z})} & P_{zz} \end{bmatrix} \begin{bmatrix} A^T \\ H^T \end{bmatrix},
\end{aligned}$$

where

$$P_{zg(\tilde{z})} = P_{zz}T^T \left(TP_{zz}T^T\right)^{-1} P_{\tilde{z}g(\tilde{z})}. \tag{13}$$

This is based on the fact that the cross term $P_{\tilde{z}g_n(\tilde{z})}$ describes the linear dependency of function $g_n(\tilde{z})$ and $\tilde{z}$ as in (9):

$$P_{\tilde{z}g_n(\tilde{z})} = P_{\tilde{z}\tilde{z}}J^T = TP_{zz}T^TJ^T. \tag{14}$$

and the term $P_{zg_n(\tilde{z})}$ is

$$P_{zg_n(\tilde{z})} = P_{zz}T^TJ^T. \tag{15}$$

Solving J from (14) and substituting it into (15) we get (13).

In the update, the matrix with cross terms is required for state variables only. This can be extracted by

$$P_{xh(z)} = \begin{bmatrix} I & 0 \end{bmatrix} P_{zh(z)}$$

when the state variables are first in the augmented state.

The algorithm for state transition is given in Algorithm 1 and an algorithm for state update is given in Algorithm 2. Use of these algorithms is beneficial when the dimension of $\tilde{z}$ is smaller than z and the matrix inverse in (13) is applied in a small dimension. Algorithms 1 and 2 are given in a general form and for applications they can be optimized further by considering the structures of matrices. Examples of exploiting the partially linear state are given in sections 7.1 and 7.2.

ALGORITHM 1: State transition using a partially linear measurement

Input: $x_{t-1} \sim \mathrm{N}(\mu_{x_{t-1}}, P_{x_{t-1}x_{t-1}})$ // State estimate from previous time step
$\varepsilon_q \sim \mathrm{N}(\mu_{\varepsilon_q}, P_{\varepsilon_q\varepsilon_q})$, // State transition noise
$P_{x_{t-1}\varepsilon_q}$ // cross covariance between state and state transition noise
$f(x, \varepsilon_q) = f(z) = Ag(Tz) + Hz$, // State transition function of given form
Output: $x \sim \mathrm{N}(\mu_{x^-}, P_{x^-x^-})$ // Propagated state

$\mu_z = \begin{bmatrix} \mu_{x_{t-1}} \\ \mu_{\varepsilon_q} \end{bmatrix}$ // Augmented mean

$P_{zz} = \begin{bmatrix} P_{x_{t-1}x_{t-1}} & P_{x_{t-1}\varepsilon_q} \\ P^T_{x_{t-1}\varepsilon_q} & P_{\varepsilon_q\varepsilon_q} \end{bmatrix}$ // Augmented covariance

$\mu_{\tilde{z}} = T\mu_z$ // Transformed mean

$P_{\tilde{z}\tilde{z}} = TP_{zz}T^T$ // Transformed covariance

Compute $\mu_{g(\tilde{z})}$,$P_{g(\tilde{z})g(\tilde{z})}$, and $P_{\tilde{z}g(\tilde{z})}$ using a KFE

$P_{zg(\tilde{z})} = P_{zz}T^T \left(TP_{zz}T^T\right)^{-1} P_{\tilde{z}g(\tilde{z})}$

$\mu_{x^-} = A\mu_{g(Tz)} + H\mu_z$ // Predicted mean of state

$P_{x^-x^-} = \begin{bmatrix} A & H \end{bmatrix} \begin{bmatrix} P_{g(\tilde{z})g(\tilde{z})} & P^T_{zg(\tilde{z})} \\ P_{zg(\tilde{z})} & P_{zz} \end{bmatrix} \begin{bmatrix} A^T \\ H^T \end{bmatrix}$ // Predicted covariance

ALGORITHM 2: Update using a partially linear measurement

Input: $x^- \sim \mathrm{N}(\mu_{x^-}, P_{x^-x^-})$ // Prior state
$\varepsilon^y \sim \mathrm{N}(\mu_{\varepsilon^y}, P_{\varepsilon^y\varepsilon^y})$, // Measurement noise
$P_{x^-\varepsilon^y}$ // cross covariance between state and measurement noise
$h(x, \varepsilon^y) = h(z) = Ag(Tz) + Hz$, // Measurement function of given form
Output: $x \sim \mathrm{N}(\mu_x, P_{xx})$ // Posterior estimate

$\mu_z = \begin{bmatrix} \mu_{x^-} \\ \mu_{\varepsilon^y} \end{bmatrix}$ // Mean of the augmented state

$P_{zz} = \begin{bmatrix} P_{x^-x^-} & P_{x^-\varepsilon^y} \\ P^T_{x^-\varepsilon^y} & P_{\varepsilon^y\varepsilon^y} \end{bmatrix}$ // Covariance matrix of the augmented state

$\mu_{\tilde{z}} = T\mu_z$ // Transformed mean

$P_{\tilde{z}\tilde{z}} = TP_{zz}T^T$ // Transformed covariance

Compute $\mu_{g(\tilde{z})}$,$P_{g(\tilde{z})g(\tilde{z})}$, and $P_{\tilde{z}g(\tilde{z})}$ using a KFE

$P_{zg(\tilde{z})} = P_{zz}T^T \left(TP_{zz}T^T\right)^{-1} P_{\tilde{z}g(\tilde{z})}$

$\mu_{h(z)} = A\mu_{g(Tz)} + H\mu_z$ // Predicted mean of measurement

$P_{h(z)h(z)} = \begin{bmatrix} A & H \end{bmatrix} \begin{bmatrix} P_{g(\tilde{z})g(\tilde{z})} & P^T_{zg(\tilde{z})} \\ P_{zg(\tilde{z})} & P_{zz} \end{bmatrix} \begin{bmatrix} A^T \\ H^T \end{bmatrix}$ // Innovation covariance

$P_{xh(z)} = \begin{bmatrix} P_{zg(\tilde{z})} & P_{zz} \end{bmatrix}_{[1:n,:]} \begin{bmatrix} A^T \\ H^T \end{bmatrix}$ // State-measurement cross correlation

$K = P_{xh(z)}P^{-1}_{h(z)h(z)}$ // Kalman Gain

$\mu_x = K(y - \mu_{h(z)})$ // Posterior mean

$P_{xx} = P_{x^-x^-} - KP_{h(z)h(z)}K^T$ // Posterior covariance

4.2 Conditionally Linear Measurements

In [Morelande and Moran 2007] a situation where a function can be divided into nonlinear z_n and conditionally linear z_l parts

$$g(z) = g_n(z_n) + G_n(z_n)z_l \tag{16}$$

is considered. In the original article the distributions of this nonlinear function is approximated using a modification of the UKF. The number of used sigma-points depends on the dimension of z_n instead of the dimension of the full state. This algorithm can be also used with other KFEs that use weighted points to compute the expectations, as in (10)-(11), but it cannot be used with other types of filters e.g. DD or SOKF2. This algorithm is converted to use the GF in [Beutler et al. 2009].

In the algorithm the sigma-points are computed for the nonlinear part only and then used for computing the conditional probabilities of the conditionally linear part:

$$\begin{aligned}\mu_{z_l|\chi_i} &= \mu_{z_l} + P_{z_l z_n} P_{z_n z_n}^{-1}(\chi_i - \mu_{z_n}) \\ P_{z_l z_l|\chi_i} &= P_{z_l z_l} - P_{z_l z_n} P_{z_n z_n}^{-1} P_{z_n z_l}.\end{aligned}$$

The matrices in the equations above are independent of the sigma-point χ_i. Thus, $P_{z_l z_n} P_{z_n z_n}^{-1}$ and $P_{z_l z_l|\chi_i}$ need to be computed only once. According to [Morelande and Moran 2007] the expectations for a function of form (16) can be approximated as

$$\begin{aligned}\mathcal{Y}_i &= g_n(\chi_i) + G_n(\chi_i)\mu_{z_l|\chi_i} \\ \mu_{g(z)} &= \sum w_i \mathcal{Y}_i \\ P_{g(z)g(z)} &= \sum w_i \left(\left(\mathcal{Y}_i - \mu_{g(z)}\right) \left(\mathcal{Y}_i - \mu_{g(z)}\right)^T + G(\chi_i) P_{z_l z_l|\chi_i} G(\chi_i)^T \right) \\ P_{z g(z)} &= \sum w_i \left(\left(\begin{bmatrix} \chi_i \\ \mu_{z_l|\chi_i} \end{bmatrix} - \mu_z \right) \left(\mathcal{Y}_i - \mu_{g(z)}\right)^T + \begin{bmatrix} 0 \\ P_{z_l z_l|z_n} G(\chi_i)^T \end{bmatrix} \right).\end{aligned}$$

These formulas can be used with algorithms 1 and 2 to solve the expectations for the nonlinear part $g(\tilde{z})$. This is done in the example in Section 7.1.

5 Optimizations based on Unobserved and Static State Variables

5.1 Part of state is unobserved and static

A state variable is considered static when its distribution does not change in the state transition and unobserved when it does not have an effect on the measurement function. This section considers the situation where some of the state variables are unobserved and static. This situation has been considered in [Davison 1998, Guivant and Nebot 2001]. The presented algorithm is based on the latter, which has more compact formulas. The noise variables are, by definition, independent of other noise variables at other times. Thus, they cannot be static and the augmented notation is not used in this section.

The state is divided into two parts, active x_a and inactive x_u. The inactive part contains variables that are static and unobserved and the active part contains the rest of the state variables.

In [Guivant and Nebot 2001] the algorithm that does updates with EKF stores the following auxiliary variables:

$$\psi_t = \psi_{t-1} + \Phi_{t-1}^T J_t^T P_{h(x_a)h(x_a),t}^{-1} J_t \Phi_{t-1} \tag{17}$$

$$c_t = c_{t-1} + \Phi_{t-1}^T J_t^T P_{h(x_a)h(x_a),t}^{-1} (y - J_t \mu_t^-) \tag{18}$$

$$\Phi_t = (I - K_t J_t) \Phi_{t-1}^T \tag{19}$$

Initial values for these variables are $\psi_{t_0} = \mathbf{0}$, $\Phi_{t_0} = I$, and $c_{t_0} = \mathbf{0}$. When needed the full estimate can be extracted by

$$\begin{aligned}
\mu_{u,t} &= \mu_{u,t_0} + P_{ua,t_0} c_t \\
P_{uu,t} &= P_{uu,t_0} - P_{au,t_0} \psi_t P_{ua,t_0} \\
P_{au,t} &= \Phi_t^T P_{au,t_0} \\
P_{ua,t} &= P_{ua,t} \Phi_t .
\end{aligned}$$

These variables contain $2n_a^2 + n_a$ elements that have to be stored (the memory requirement can be reduced by noting that ψ_t is symmetric), but only the matrix of the latest time step is needed. When $n_a \ll n_u$ this is negligible. It is worthwhile to note that only term that contains a matrix inverse $J_t^T P_{h(x_a)h(x_a),t}^{-1}$ is used in equations (17) and (18) and also in the computation of the Kalman gain for EKF and for speeding up the algorithm it is required that it is computed only once.

Using (9) we can rewrite (17)-(19) in a more general form

$$\begin{aligned}
\psi_t &= \psi_{t-1} + \Phi_{t-1}^T P_{aa}^{-1} P_{ah(x)} P_{h(x_a)h(x_a),t}^{-1} P_{h(x)a} P_{aa}^{-1} \Phi_{t-1} \\
c_t &= c_{t-1} + \Phi_{t-1}^T P_{aa}^{-1} P_{ah(x)} P_{h(x_a)h(x_a),t}^{-1} (y - \mu_t^-) \\
\Phi_t &= (I - K_t P_{ah(x)}^T P_{aa}^{-1}) \Phi_{t-1} .
\end{aligned}$$

This formulation can be used with a wider set of KFEs than previous, but it contains the inverse of the active state P_{aa}^{-1}. The computation of this inverse is feasible when the active state is small. The full algorithm is given in Algorithm 3.

5.2 Part of state is unobserved and whole state is static

In [Lefebvre et al. 2005], an update algorithm for the situation when all state variables are static is presented. In this algorithm the update is done only for observed variables at every time step. When the full state estimate is required at time index $t_0 + \Delta$ it can be extracted by:

$$L = P_{ua,t_0} P_{aa,t_0}^{-1} \tag{20}$$

$$\begin{aligned}
\mu_{u,t_0+\Delta} &= \mu_{u,t_0} + L\left(\mu_{a,t_0+\Delta} - \mu_{a,t_0}\right) \\
P_{ua,t_0+\Delta} &= L P_{aa,t_0+\Delta}
\end{aligned}$$

$$P_{uu,t_0+\Delta} = P_{uu,t_0} - L\left(P_{aa,t_0} - P_{aa,t_0+\Delta}\right) L^T . \tag{21}$$

This algorithm does not require any extra variables stored during the estimation process. An example of the use of this optimization is given in Section 7.3.

ALGORITHM 3: State estimation when a part of the state is unobserved and static.

Input: $x_{t_0} = \begin{bmatrix} x_{a,t_0} \\ x_{u,t_0} \end{bmatrix} \sim \mathrm{N}\left(\begin{bmatrix} \mu_{a,t_0} \\ \mu_{u,t_0} \end{bmatrix}, \begin{bmatrix} P_{aa,t_0} & P_{au,t_0} \\ P_{ua,t_0} & P_{uu,t_0} \end{bmatrix} \right)$ `// Initial state`

$x_t = h_t(x_{a,t}, \varepsilon_t^y)$, `// Measurement model`

$x_{t+1} = \begin{bmatrix} x_{a,t+1} \\ x_{u,t+1} \end{bmatrix} = \begin{bmatrix} f_{t+1}(x_{a,t}, \varepsilon^x) \\ x_{u,t} \end{bmatrix}$ `// State transition model`

Output: $x_{t_{0+\Delta}} = \begin{bmatrix} x_{a,t_{0+\Delta}} \\ x_{u,t_{0+\Delta}} \end{bmatrix} \sim \mathrm{N}\left(\begin{bmatrix} \mu_{a,t_{0+\Delta}} \\ \mu_{u,t_{0+\Delta}} \end{bmatrix}, \begin{bmatrix} P_{aa,t_{0+\Delta}} & P_{au,t_{0+\Delta}} \\ P_{ua,t_{0+\Delta}} & P_{uu,t_{0+\Delta}} \end{bmatrix} \right)$ `// State at` $t_{0+\Delta}$

Initialize auxiliary variables:

$\psi_{t_0} = \mathbf{0}$

$\Phi_{t_0} = I$

$c_{t_0} = \mathbf{0}$

for $t = t_0 + 1$ *to* $t_{0+\Delta}$ **do**

Prediction of the active part:

$\mu_{a,t}^- = \mu_{f_t(x_{a,t-1}, \varepsilon_t^x)}$ `// Predicted mean of active part of the state`

$P_{aa,t}^- = P_{f(x_{a,t-1}\varepsilon_t^x), f(x_{a,t-1}, \varepsilon_t^x)}$ `// Predicted covariance of the active part`

Update of the active part:

$y_t^- = \mu_{h_t(x_{a,t}^-, \varepsilon_t^y)}$

$S_t = P_{h_t(x_{a,t}^-, \varepsilon_t^y) h_t(x_{a,t}^-, \varepsilon_t^y)}$

$K_t = P_{x_a h_t(x_{a,t}^-, \varepsilon_t^y)} S_i^{-1}$

$\mu_{a,t} = \mu_{a,t}^- + K(y_t - y_t^-)$

$P_{aa,t} = P_{aa,t}^- - K_t S_t K_t^T$

Update auxiliary variables:

$\psi_t = \psi_{t-1} + \Phi_{t-1}^T P_{aa,t}^{-1} P_{ah(x),t} P_{h_t(x_{a,t}) h_t(x_a),t}^{-1} P_{h(x_{a,t}) a_t} P_{aa,t}^{-1} \Phi_{t-1}$

$c_t = c_{t-1} + \Phi_{t-1}^T P_{aa,t}^{-1} P_{ah(x),t} P_{h(x_a) h(x_a),t}^{-1} (y_t - y_t^-)$

$\Phi_t = (I - K_t P_{ah(x),t}^T P_{aa,t}^{-1}) \Phi_{t-1}$

end

Compute the full mean and covariance:

$\mu_{x,t_{0+\Delta}} = \begin{bmatrix} \mu_{a,t_0+\Delta} \\ \mu_{u,t_0} + P_{ua,t_0} c_{t_{0+\Delta}} \end{bmatrix}$ `// Posterior mean`

$P_{xx,t_{0+\Delta}} = \begin{bmatrix} P_{aa,t_0+\Delta} & \Phi_{t_{0+\Delta}}^T P_{au,t_0} \\ \left(\Phi_{t_{0+\Delta}}^T P_{au,t_0}\right)^T & P_{uu,t_0} - P_{au,t_0} \psi_{t_{0+\Delta}} P_{ua,t_0} \end{bmatrix}$ `// Posterior covariance`

ALGORITHM 4: Block Update KF

Input: $x_0 \sim \mathrm{N}(\mu_{x,0}, P_{xx,0})$ `// Prior state`
$\varepsilon_i^y \sim \mathrm{N}(\mu_{\varepsilon_i^y}, P_{\varepsilon_i^y \varepsilon_i^y})$, $1 \leq i \leq n$ `// Measurement noises`
$h_i(x, \varepsilon_i^y)$, $1 \leq i \leq n$ `// Measurement functions`
Output: $x_n \sim \mathrm{N}(\mu_{x,n},, P_{xx,n})$ `// Posterior estimate`
for *i=1 to n* **do**
$y_i^- = \mu_{h_i(x_{i-1}, \varepsilon_i^y)}$
$S_i = P_{h_i(z_{t,i}^-) h(x_{i-1}, \varepsilon_i^y)}$
$K_i = P_{x h_i(z_{t,i}^-)} S_i^{-1}$
$\mu_{x,i} = \mu_{x,i-1} + K(y_i - y_i^-)$
$P_{xx,i} = P_{xx,i}^- - K_i S_i K_i^T$
end

6 Optimizations related to the inverse of the innovation covariance

6.1 Block diagonal measurement covariance

The equation for the Kalman gain (2) contains the inverse of the innovation covariance S. When the dimension of the innovation covariance (i.e. the number of measurements) is large the computation of the inverse or solving the set of linear equations is a computationally expensive operation.

When measurements are linear and the measurement covariance is diagonal, the Kalman update can be applied one measurement element at a time and the partially updated state is used as a prior for the next measurement [Stengel 1986]. This kind of update can be generalized also to block diagonal measurement covariances and the update can be done by applying one block at the time.

This reduces the computation time when the state dimension is small, the measurement dimension is large, and the measurements are independent. When some measurements are independent, the KFE has superlinear complexity and different measurement elements contain different state terms this can be effectively combined with the algorithms presented in Section 4. The block update does not change the estimate when the measurement model is linear, but when the model is nonlinear the linearization is redone in the partially updated state. This may alter the output of the algorithm. In [Raitoharju et al. 2015] this is used to improve the estimation accuracy by applying the most linear measurements first. The algorithm for updating by blocks is given in Algorithm 4.

6.2 Applying the matrix inversion lemma to innovation covariance

When the measurements are not independent the block update formula cannot be used. In some situations the innovation covariance can be written in the form:

$$S = P_s + U P_v U^T, \tag{22}$$

where P_s is easy to invert, P_s has a small dimension and U is a transformation matrix.

Using matrix inversion lemma (a.k.a. the Woodbury formula, see, for example, [Henderson and Searle 1981]) the inverse of (22) is

$$S^{-1} = P_s^{-1} - P_s^{-1} U \left(P_v^{-1} + U P_s^{-1} U^T \right)^{-1} U^T P_s^{-1}. \tag{23}$$

This formula is worth using if the inverse of P_s is easy to compute or can be computed offline and the dimension of P_v is smaller than the dimension of the P_s. An example of its use is given in Section 7.2.

7 Example Applications

7.1 Pedestrian Dead Reckoning

In Pedestrian Dead Reckoning (PDR) a pedestrian's steps and heading changes are measured with accelerometers and gyroscopes. The change of position is based on these. In [Nurminen et al. 2013] a state consists of two-dimensional position r, heading θ, step length l, and possibly a floor index. The state transition is computed with a particle filter. Here we show how the state model without floor index and with normal assumed noise can be optimized for use with KFEs. We use in this example algorithms presented in Sections 4.1 and 4.2.

The nonlinear state transition model is

$$x_{t+1} = \begin{bmatrix} r_{1,t+1} \\ r_{2,t+1} \\ \theta_{t+1} \\ l_{t+1} \end{bmatrix} = \begin{bmatrix} r_{1,t} + (l_t + \varepsilon_l) \cos(\theta_t + \varepsilon_\theta) + \varepsilon_{r,1} \\ r_{2,t} + (l_t + \varepsilon_l) \sin(\theta_t + \varepsilon_\theta) + \varepsilon_{r,2} \\ \theta_t + \varepsilon_\theta \\ l_t + \varepsilon_l \end{bmatrix}, \tag{24}$$

where $\varepsilon_\theta \sim \mathrm{N}(\Delta_t, \sigma_\theta^2)$ is the noisy heading change, with Δ_t measured from gyroscopes, l is the footstep length and r is the position of the pedestrian. The augmented state is $z = \begin{bmatrix} r_1 & r_2 & \theta & l & \varepsilon_{r,1} & \varepsilon_{r,2} & \varepsilon_\theta & \varepsilon_l \end{bmatrix}^T$. The state transition model can be written now in the form of (12):

$$\begin{aligned} \tilde{z} &= \begin{bmatrix} \theta_t + \varepsilon_\theta \\ l_t + \varepsilon_l \end{bmatrix} \\ T &= \begin{bmatrix} \mathbf{0}_{2\times 2} & I_{2\times 2} & \mathbf{0}_{2\times 2} & I_{2\times 2} \end{bmatrix} z \\ g(\tilde{z}) &= \begin{bmatrix} \tilde{z}_2 \cos \tilde{z}_1 \\ \tilde{z}_2 \sin \tilde{z}_1 \end{bmatrix} \\ A &= \begin{bmatrix} I_{2\times 2} \\ \mathbf{0}_{2\times 2} \end{bmatrix} \\ H &= \begin{bmatrix} I_{4\times 4} & I_{4\times 4} \end{bmatrix} \end{aligned} \tag{25}$$

In this formulation the reduced state $\tilde{z}$ is two-dimensional and the nonlinear part of the state transition function is also two-dimensional.

If this problem is to be solved with a KFE that uses sigma-points, we can optimize the implementation further. This can be done because $\tilde{z}_2$ in (25) is conditionally linear, given $\tilde{z}_1$. The parts of function (16) are

$$\begin{aligned} z_n &= \tilde{z}_1 \\ z_l &= \tilde{z}_2 \end{aligned}$$

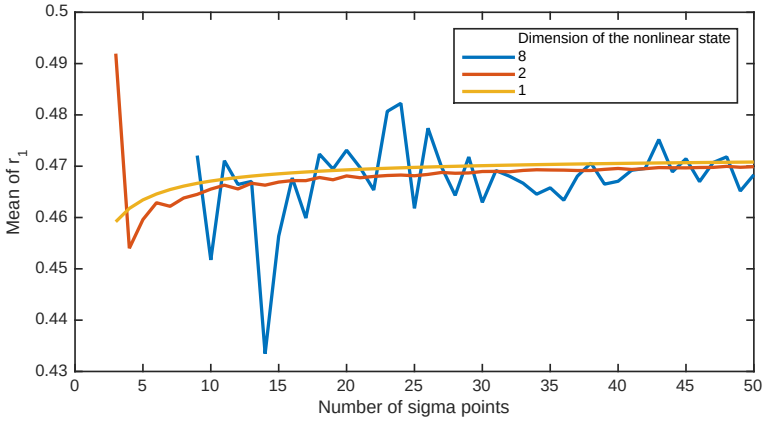

Figure 1: Estimated mean of r_1 with different number of function evaluation points using S^2KF

$$g_n(\tilde{z}_1) = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$
$$G_n(\tilde{z}_1) = \begin{bmatrix} \sin \tilde{z}_1 \\ \cos \tilde{z}_1 \end{bmatrix}$$

Using this the dimension of the nonlinear part of the state is reduced to 1 and the sigma-points are generated for only one dimension.

Figure 1 shows mean estimates of the first position variable (r_1) after one propagation step in a test. The estimates are computed using asymmetric sigma-points that are based on Gaussian localized cumulative disrtibution. Implementation of the sigma-point generator and S^2KF can be found in [Steinbring]. S^2KF is applied to original state (24) and to above presented reduced states with 2 and 1 nonlinear variables. The propagation of the state with the full state model cannot be computed with fewer than 9 sigma-points. The state computed using the model that uses conditional linearity in addition to the linearity has the smoothest convergence to an estimate and the estimate with full state estimate has the most variance. This shows that the use of optimizations can also improve the estimation accuracy.

7.2 Source tracking using a microphone array Time Difference of Arrival (TDoA)

In tracking of a sound source using a microphone array the received acoustic signals of two microphones are compared so that the TDoA of the acoustic signal is measured. In ideal situation the TDoA is

$$h_{i,j}(r_0) = \frac{\|r_i - r_0\| - \|r_j - r_0\|}{v},$$

where r_0 is the location of the sound source, r_i is the location of ith microphone and v is the speed of sound. When there are m microphones in the array there are at most $\frac{m(m-1)}{2}$ TDoA measurements available [Oualil et al. 2012].

In practice the measurements contain noise. Here we consider model with noises

$$h_{i,j}(x) = \|r_i - x\| - \|r_j - x\| + \varepsilon_i - \varepsilon_j + \varepsilon_{i,j},$$

where ε_i is the noise corresponding to ith microphone and $\varepsilon_{i,j}$ is error corresponding to the correlation of the measured signals from the microphone pair and the speed of sound is multiplied out from the equation.

When $\varepsilon_{i,j} = 0$ for all microphone pairs the measurement noises are not independent for all possible pairs and the measurements can be modeled with $m-1$ measurements of form

$$h_i(x) = \|r_i - x\| - \|r_m - x\| + \varepsilon_i - \varepsilon_m.$$

This kind of assumption is done for example in [Bechler et al. 2004]. In practice, this assumption does not hold and the selection of microphone pairs is a tradeoff, where pairs close to each other have smaller correlation error $\varepsilon_{i,j}$, but worse geometry than pairs far away from each other [Silverman et al. 1998]. Here we consider the situation where all possible microphone pairs are used and errors ε_i and $\varepsilon_{i,j}$ are modeled as Gaussians. The augmented state model is

$$\begin{bmatrix} x^T & \varepsilon_{1,1} & \varepsilon_{1,2} & \dots & \varepsilon_{m,m} & \varepsilon_1 & \dots & \varepsilon_m \end{bmatrix}^T,$$

where x contains 3 position and velocity variables.

Using (10)-(11) the nonlinear part of the estimation can be done only for the state variables and the sigma-points would be required only for 6 dimensional state instead of using sigma-points also for $\frac{m(m+1)}{2}$ noise terms. The measurement model can be written in compact form using (12):

$$\begin{aligned}
\tilde{z} &= x_{1:3} \\
T &= \begin{bmatrix} I_{3\times 3} & \mathbf{0_{3\times 3}} \end{bmatrix} \\
g(\tilde{z}) &= \begin{bmatrix} \|r_1 - \tilde{z}\| \\ \|r_2 - \tilde{z}\| \\ \vdots \\ \|r_m - \tilde{z}\| \end{bmatrix} \\
A &= \begin{bmatrix} \mathbf{1}_{m-1\times 1} & -I_{m-1\times m-1} & \\ \mathbf{0}_{m-2\times 1} & \mathbf{1}_{m-2\times 1} & -I_{m-2\times m-2} \\ \mathbf{0}_{m-3\times 2} & \mathbf{1}_{m-3\times 1} & -I_{m-3\times m-3} \\ & \vdots & \\ \mathbf{0}_{1\times m-2} & 1 & -1 \end{bmatrix} \\
H &= \begin{bmatrix} \mathbf{0}_{\frac{m(m-1)}{2}\times 6} & A & I_{\frac{m(m-1)}{2}\times\frac{m(m-1)}{2}} \end{bmatrix}.
\end{aligned}$$

Using these the dimension of the nonlinear function is reduced from $\frac{m(m-1)}{2}$ to m. This reduces the number of updated elements for each sigma-point in (7) from $\left(\frac{m(m-1)}{2}\right)^2$ to m^2. A and H matrices are sparse and an application of sparse linear algebra codes would further enhance the performance of the algorithm.

The dimension of S is $\left(\frac{m(m-1)}{2}\right)^2 \times \left(\frac{m(m-1)}{2}\right)^2$. The computation of the inverse of the innovation covariance for Kalman gain (2) can be optimized using

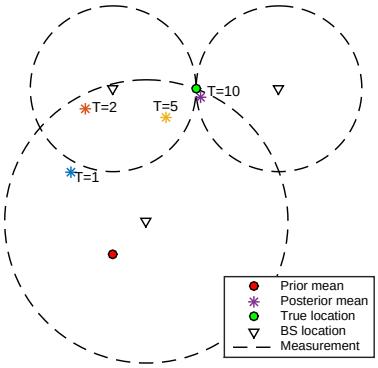

Figure 2: Posterior means computed with RUF with varying number of iterations.

the inversion formula (23). The noise terms are assumed independent and their covariance matrices are diagonal. We split noises into two parts, whose covariance matrices are D_1 and D_2. D_1 corresponds to microphone specific terms ε_i^y and D_2 to microphone pair specific terms $\varepsilon_{i,j}^y$. Because $\tilde{z}$ contains only state variables and they do not contribute to the linear part ($AP_{zg(\tilde{z})^T}H^T = \mathbf{0}$), the innovation covariance can be written as

$$S = D_2 + A(P_{g(\tilde{z})g(\tilde{z})} + D_1)A^T$$

and its inverse is

$$S^{-1} = D_2^{-1} + D_2^{-1}A\left[(P_{g(\tilde{z})g(\tilde{z})} + D_1)^{-1} + AD_2^{-1}A^T\right]^{-1}A^T D_2^{-1}.$$

Because D_2^{-1} is diagonal its inverse is diagonal with the reciprocal of the diagonal elements of D_2 on its diagonal. Other inverses that are in this equations are applied to $m \times m$ matrices instead of $\left(\frac{m(m-1)}{2}\right)^2 \times \left(\frac{m(m-1)}{2}\right)^2$ matrices.

7.3 Optimization of iterative KFEs

Optimizations presented in sections 5.1 and 5.2 are most widely used in the field of Simultaneous Localization and Mapping (SLAM). In SLAM, the state contains dynamic part that represents the agent (robot) that is mapping the environment and localizing itself. The environment is mapped using static landmarks.

We propose that these can be used also to optimize KFEs that do the update iteratively e.g. Recursive Update Filter (RUF) [Zanetti 2012] and its extension for sigma-point filters [Huang et al. 2015] or posterior linearization filters [García-Fernández et al. 2015]. In RUF, the state is updated by applying the update in parts that have smaller impact than the original update. After each part of the update the linearization is recomputed. These iterative filters are developed to that they assume additive Gaussian noise. Instead of making the updates for full state the update could be done iteratively only for the observed variables. Because the iterative update is actually one update, the unobserved state variables are static during the partial updates.

Table 1: Summary of presented optimizations

Section	Exploited structure
Section 4.1	Partially linear functions $f(z) = Ag(Tz) + Hz$
Section 4.2	Conditionally linear functions $f(z) = g_n(z_n) + G(z_n)z_l$
Section 5.1	Part of state is unobserved and static
Section 5.2	Part of state is unobserved and whole state is static
Section 6.1	Measurement covariance is block diagonal
Section 6.2	Innovation covariance of form $S = D_2 + A(P_{g(\bar{z})g(\bar{z})} + D_1)A^T$

In this example we consider a state model with 3D position, velocity, and acceleration i.e. 9 state variables. Measurements used are range measurements from 3 Basestations (BSs). The measurement model for range from the ith BS is

$$y_i = \left\| x_{[1:3]} - r_i \right\| + \varepsilon_i^y, \tag{26}$$

where r_i is the location of the ith BS. In our example, the prior has a large variance and its mean is located so that the linearization in the mean is not accurate. Figure 2 shows the example situation and the posterior means which are computed with different number of iterations of RUF. The estimate with one iteration ($T = 1$) is identical to EKF estimate and is not close to the true location. The estimate with 10 iterations is close to the true location. The computation of 10 iterations with RUF involves the updates to the 9×9 state covariance matrices 10 times. Because the measurement model (26) depends only on 3 first state variables the active state can be selected to consist only those and the iterative update is computed only for the position variables. The full covariance update is computed after all iterations are applied. In this update scheme, the 3×3 covariance matrix is updated T times and the 9×9 matrix is updated only once using (20)-(21).

8 Conclusions

In this survey, we collected different optimizations for KFEs that exploit the structure of the state and measurement models to reduce the computational load. Table 1 summarizes the structures of models that are exploited in different optimizations. These structures are present in various nonlinear estimation problems.

Table 2: Summary of properties of selected KFEs

Algorithm	Order of measurement function evaluations. α = algorithm specific parameter ($\alpha > 0$), j = number of iterations	Computation of J A=analytically, I=using (9), N=numerically inside algorithm	Expectations of form (6)-(8)	Iterative algorithm
Extended Kalman Filter (EKF) [Jazwinski 1970]	1	A		
Second Order Extended Kalman Filter (EKF2) [Jazwinski 1970]	1	A		
Unscented Kalman Filter (UKF) [Wan and Van der Merwe 2000]	n	I	X	
Cubature Kalman Filter (CKF) [Arasaratnam and Haykin 2009]	n^α	I	X	
Smart Sampling Kalman Filter (S^2KF) [Steinbring and Hanebeck 2014]	$n + \alpha$	I	X	
Gaussian Filter (GF) [Steinbring and Hanebeck 2014]	αn	I	X	
Gauss-Hermite Quadrature Kalman filter (QKF) [Ito and Xiong 2000]	α^n	I	X	
Divided Difference (DD) [Nørgaard et al. 2000]	n	N		
Second Order Polynomial Kalman Filter (SOKF2) [Ito and Xiong 2000]	n^2	N		
Recursive Update Filter (RUF) [Zanetti 2012]	j	A		X
Unscented Recursive Update Filter (URUF) [Huang et al. 2015]	jn	I	X	X

Table 2 shows some properties of selected KFEs. Different optimizations give different amount of speed increase with different KFEs. Optimizations presented in sections 4.1 and 4.2 are the most useful with KFEs that have a high number of function evaluations as a function of state dimension. The exploitation of conditionally linear part of functions in Section 4.2 requires that the expectations are approximated with equations of form (6)-(8). The optimization that exploits the partially unobserved and static state in Section 5.1 is better suited for algorithms that compute the J within the algorithm (A or N in Table 2). In the example in Section 7.3 we showed that the optimization that exploits static state and partly unobserved variables is beneficial with iterative KFEs.

In addition to gaining increase in speed of computation the algorithms may also lead to better estimation accuracy. An example of this was given in Section 7.1.

References

Ienkaran Arasaratnam and Simon Haykin. Cubature Kalman filters. *IEEE Transactions on Automatic Control*, 54(6):1254–1269, June 2009. ISSN 0018-9286. doi: 10.1109/TAC.2009.2019800.

Yaakov Bar-Shalom, Thiagalingam Kirubarajan, and X.-Rong Li. *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, Inc., New York, NY, USA, 2002. ISBN 0471221279.

Dirk Bechler, Markus S Schlosser, and Kristian Kroschel. System for robust 3d speaker tracking using microphone array measurements. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 3, pages 2117–2122, Sendai, Japan, 2004. IEEE. doi: 10.1109/IROS.2004.1389722.

Frederik Beutler, Marco F. Huber, and Uwe D. Hanebeck. Gaussian filtering using state decomposition methods. In *Proceedings of the 12th International Conference on Information Fusion (FUSION)*, pages 579–586, Seattle, WA, USA, July 2009. IEEE.

Mark Briers, Simon R. Maskell, and Robert Wright. A Rao-Blackwellised Unscented Kalman filter. In *Proceedings of the 6th International Conference of Information Fusion (FUSION)*, volume 1, pages 55–61, Cairns, Australia, July 2003. IEEE. doi: 10.1109/ICIF.2003.177426.

Andrew John Davison. *Mobile Robot Navigation Using Active Vision*. PhD thesis, Department of Engineering Science, University of Oxford, 1998.

Arnaud Doucet, Nando de Freitas, Kevin P. Murphy, and Stuart J. Russell. Rao-Blackwellised particle filtering for dynamic Bayesian networks. In *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*, UAI '00, pages 176–183, San Francisco, CA, USA, 2000. Morgan Kaufmann Publishers Inc. ISBN 1-55860-709-9. URL `http://dl.acm.org/citation.cfm?id=647234.720075`.

Ángel F. García-Fernández, Lennart Svensson, Mark R. Morelande, and Simo Särkka. Posterior linearisation filter: principles and implementation using sigma points. *IEEE Transactions on Signal Processing*, PP(99):1–1, 2015. ISSN 1053-587X. doi: 10.1109/TSP.2015.2454485.

José E. Guivant and Eduardo Mario Nebot. Optimization of the simultaneous localization and map-building algorithm for real-time implementation. *IEEE Transactions on Robotics and Automation*, 17(3):242–257, Jun 2001. ISSN 1042-296X. doi: 10.1109/70.938382.

H. V. Henderson and S. R. Searle. On deriving the inverse of a sum of matrices. *SIAM Review*, 23(1):pp. 53–60, 1981. ISSN 00361445. URL `http://www.jstor.org/stable/2029838`.

Petri Hiltunen, Simo Särkkä, Ilkka Nissilä, Atte Lajunen, and Jouko Lampinen. State space regularization in the nonstationary inverse problem for diffuse optical tomography. *Inverse Problems*, 27(2):025009, 2011.

Yu-Chi Ho and Robert Lee. A Bayesian approach to problems in stochastic estimation and control. *IEEE Transactions on Automatic Control*, 9(4):333 – 339, October 1964. doi: 10.1109/TAC.1964.1105763.

Yulong Huang, Yonggang Zhang, Ning Li, and Lin Zhao. Design of sigma-point kalman filter with recursive updated measurement. *Circuits, Systems, and Signal Processing*, pages 1–16, 2015. ISSN 0278-081X. doi: 10.1007/s00034-015-0137-y.

Kazufumi Ito and Kaiqi Xiong. Gaussian filters for nonlinear filtering problems. *IEEE Transactions on Automatic Control*, 45(5):910–927, May 2000. ISSN 0018-9286. doi: 10.1109/9.855552.

Andrew H. Jazwinski. *Stochastic Processes and Filtering Theory*, volume 64 of *Mathematics in Science and Engineering*. Academic Press, New York, NY, USA, 1970.

Matthew James Keeling and Pejman Rohani. *Modeling infectious diseases in humans and animals*. Princeton University Press, Princeton, 2008. ISBN 978-0-691-11617-4. URL `http://opac.inria.fr/record=b1124487`.

Tine Lefebvre, Herman Bruyninckx, and Joris De Schutter. *Nonlinear Kalman Filtering for Force-Controlled Robot Tasks*, volume 19 of *Springer Tracts in Advanced Robotics*. Springer-Verlag, Berlin, Germany, 2005. ISBN 978-3-540-28023-1.

Mark R. Morelande and Bill Moran. An unscented transformation for conditionally linear models. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 3, pages III–1417–III–1420, Honolulu, HI, USA, April 2007. IEEE. doi: 10.1109/ICASSP.2007.367112.

Mark R. Morelande and Branko Ristic. Reduced sigma point filtering for partially linear models. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 3, pages III–III, Toulouse, France, May 2006. IEEE. doi: 10.1109/ICASSP.2006.1660584.

Philipp Müller, Henk Wymeersch, and Robert Piché. UWB positioning with generalized Gaussian mixture filters. *IEEE Transactions on Mobile Computing*, 13 (10):2406–2414, Oct 2014. ISSN 1536-1233. doi: 10.1109/TMC.2014.2307301.

Magnus Nørgaard, Niels K. Poulsen, and Ole Ravn. New developments in state estimation for nonlinear systems. *Automatica*, 36(11):1627–1638, November 2000. ISSN 0005-1098. doi: 10.1016/S0005-1098(00)00089-3.

Henri Nurminen, Anssi Ristimaki, Simo Ali-Löytty, and Robert Piché. Particle filter and smoother for indoor localization. In *Proceedings of the International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pages 1–10, Montbéliard, France, Oct 2013. IEEE. doi: 10.1109/IPIN.2013.6817903.

Youssef Oualil, Friedrich Faubel, Mathew Magimai Doss, and Dietrich Klakow. A TDOA Gaussian mixture model for improving acoustic source tracking. In *Proceedings of the 20th European Signal Processing Conference (EUSIPCO)*, pages 1339–1343, Bucharest, Romania, 2012. IEEE.

Matti Raitoharju, Robert Piché, Juha Ala-Luhtala, and Simo Ali-Löytty. Partitioned update kalman filter. *Journal of Advances in Information Fusion*, 2015. URL `http://arxiv.org/abs/1503.02857`. In press.

Simo Särkkä. *Bayesian filtering and smoothing*, volume 3. Cambridge University Press, Cambridge, UK, 2013.

Harvey F Silverman, William RIII Patterson, and James L Flanagan. The huge microphone array. *IEEE Concurrency*, 6(4):36–46, 1998. doi: 10.1109/4434.736423.

Harold W. Sorenson and Daniel L. Alspach. Recursive Bayesian estimation using Gaussian sums. *Automatica*, 7(4):465–479, 1971. doi: 10.1016/0005-1098(71)90097-5.

Jannik Steinbring. Nonlinear estimation toolbox. URL `https://bitbucket.org/nonlinearestimation/toolbox`.

Jannik Steinbring and Uwe D. Hanebeck. LRKF Revisited: The Smart Sampling Kalman Filter (SKF). *Journal of Advances in Information Fusion*, 9(2):106 – 123, December 2014. URL `http://isif.org/node/189`.

Robert F. Stengel. *Stochastic optimal control*. Wiley-Interscience, New York, 1986.

Eric A. Wan and Rudolph Van der Merwe. The unscented Kalman filter for nonlinear estimation. In *Proceedings of the Adaptive Systems for Signal Processing, Communications, and Control Symposium (AS-SPCC)*, pages 153–158, Lake Louise, AB, Canada, 2000. IEEE. doi: 10.1109/ASSPCC.2000.882463.

Renato Zanetti. Recursive update filtering for nonlinear estimation. *IEEE Transactions on Automatic Control*, 57(6):1481–1490, June 2012. ISSN 0018-9286. doi: 10.1109/TAC.2011.2178334.